arXiv:1704.00963v1 [stat.ML] 4 Apr 2017

Bayesian optimization with virtual derivative sign observations

Eero Siivola[1], **Aki Vehtari**[1], **Jarno Vanhatalo**[2], **Javier González**[3*]
[1]Aalto University, Dept. of Computer Science, {eero.siivola, aki.vehtari}@aalto.fi
[2]University of Helsinki, Dept. of Math. and Stat., and Dept. of Biosciences, jarno.vanhatalo@helsinki.fi
[3]Amazon.com, gojav@amazon.com

Abstract

Bayesian optimization (BO) is a global optimization strategy designed to find the minimum of expensive black-box functions g typically defined on a continuous sets of $\mathcal{R}^d$. Using a Gaussian process (GP) as a surrogate model for the objective and an acquisition function to systematically search its domain, BO strategies aim to minimize the amount of samples required to find the minimum of g. Although currently available acquisition functions address this goal with different degree of success, an over-exploration effect of the contour of g is typically observed. This is due to the myopic nature of most acquisitions that greedily try to over-reduce uncertainty in the border of the search domain. In most real problems, however, like the configuration of machine learning algorithms, the function domain is conservatively large and with a high probability the global minimum is not at the boundary. We propose a method to incorporate this knowledge into the searching process by adding virtual derivative observations at the borders of the search space. We use the properties of GP models that allow us to easily impose conditions on the partial derivatives of the objective. The method is applicable with any acquisition function, it is easy to use and consistently reduces the number of evaluations required to find the minimum of g irrespective of the acquisition used. We illustrate the benefits our approach in a simulation study with a battery of objective functions.

1 INTRODUCTION

Global optimization is a common problem in a very broad range of applications. Formally, it is defined as finding $\mathbf{x}_{\text{min}} \in \mathcal{X} \subset \mathcal{R}^d$ such that

$$\mathbf{x}_{\text{min}} = \arg\min_{\mathbf{x}\in\mathcal{X}} g(\mathbf{x}),$$

where $\mathcal{X}$ is generally considered to be a bounded set. In this work, we focus on cases in which g is a black-box function whose explicit form is unknown and that it is expensive to evaluate. This implies that we need to find the $\mathbf{x}_{\text{min}}$ with a finite, typically, small number of evaluations, which transform the original *optimization* problem in a sequence of *decision* problems.

Bayesian optimization of black-box functions using Gaussian Processes (GPs) as surrogate priors has become popular in recent years (see, e.g., review by Shahriari et al., 2016b). It has been proven to be an efficient method to treat the decision problems of where to evaluate g using statistical *inference*. This is typically done by means of an acquisition function that change after each point is collected and that balances and *exploration* and *exploration* in the search.

A common problem that has not been systematically studied in the BO literature is the tendency of most acquisitions to over-explore the boundary of the function domain $\mathcal{X}$. This issue is not relevant if the global minimum may lie on the border of the search space but in most cases, including when the searching is done in ubounded (Shahriari et al., 2016a), this is not the case. This effect has also been observed in the active learning literature (Krause and Guestrin, 2007) and it is known to appear when the search is done myopically, as it is the case in most acquisitions functions. Non-myopic approaches in BO (González et al., 2016; Osborne et al., 2009) can potentially deal with this problem but they are typically very expensive to compute.

In this paper we propose a new approach to correct the

[*] Work done while Javier González was at the University of Sheffield

boundary over-exploration effect of most acquisitions without having to assume the computational overhead of currently available non-myopic methods. When it is known that the global minimum does not lie on the boundary of $\mathcal{X}$, this information can be embedded in the problem: it holds that the gradient of the underlying function points towards the centre on all borders. In practice, this property can be used by including additional prior information about the shape of the function by means of virtual derivative observations. In other words, we use non-observed partial derivative data to support our prior assumption about the shape of the underlying function (its gradients on the boundary of the domain). As the derivative of a GP is also a GP (O'Hagan, 1992), including virtual derivative observations in GPs is feasible. As shown in this work, this reduces the number of required function evaluations giving rise to a battery of more efficient BO methods.

Derivative observations have been used before in the BO and GP context. Koistinen et al. (2016) use Gaussian processes with derivative observations for Bayesian optimization (BO) of minimum energy path transitions of atomic rearrangements and Wu et al. (2017) use derivative observations to decrease the number of observations needed for finding the function optimum.

Derivative observations can be used to provide shape priors. To constrain a function to have a mode in a specified location Gosling et al. (2007) add a virtual observations of first derivative being zero and the second derivative being negative. Riihimäki and Vehtari (2010) use virtual derivative observations, where only the sign of the derivative is known, to add monotonicity information. Michael and Víctor (2016) use in similar way virtual derivative observations for convexity and quasi-convexity constraints. To handle inference for the non-Gaussian contribution of the derivative sign information, Gosling et al. (2007), Riihimäki and Vehtari (2014), and Michael and Víctor (2016) use rejection sampling, Riihimäki and Vehtari (2010) use expectation propagation (EP), and Wang and Berger (2016) use Markov chain Monte Carlo (MCMC). Unlike rejection and projection based shape constraints, monotonicity via virtual observations scales moderately well with a number of dimensions as shown by Riihimäki and Vehtari (2010) and Siivola et al. (2016).

This paper is structured as follows. Section 2 briefly presents the background of GPs and BO. In section 3, we present the proposed method for including the prior of no minima being on the edges to the GP and BO. Section 4 presents the experiments that were performed with the proposed method. Finally, section 5 concludes the paper and identifies possible topics for further research.

2 THEORETICAL BACKGROUND

In this section we briefly revise theoretical background behind GPs and BO.

2.1 GAUSSIAN PROCESSES WITH PARTIAL DERIVATIVE OBSERVATIONS

Let $\mathbf{x}$ be a d-dimensional input vector. The GP prior is directly specified on the latent function $f(\mathbf{x})$ and the prior assumptions are encoded in the covariance function $k(\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$, which specifies the covariance of two latent function values $f(\mathbf{x}^{(1)})$ and $f(\mathbf{x}^{(2)})$. A zero mean Gaussian process prior can mathematically be formulated as,

$$p(\mathbf{f}) = \mathrm{N}(\mathbf{f}|\mathbf{0}, \mathbf{K}), \tag{1}$$

where $\mathbf{K}$ is a covariance matrix between n latent values $\mathbf{f}$ at input used for training, $\mathbf{X} = \left(\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(n)}\right)$, s.t. $\mathbf{K}_{ij} = k(\mathbf{x}^{(i)}, \mathbf{x}^{(j)})$.

In regression, n noisy observations $\mathbf{y}$ and o latent function values $\mathbf{f}_*$ at the test inputs $\mathbf{X}_*$ are assumed to have Gaussian relationship. With the noise variance σ^2, the covariance between the latent values at the training and test inputs $\mathbf{K}_*$, the covariance matrix of the latent values at the test inputs $\mathbf{K}_{**}$ and n dimensional identity matrix $\mathbf{I}$, the joint distribution of the observations and latent values at the test inputs is

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{f}_* \end{bmatrix} \sim \mathrm{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K} + \sigma^2\mathbf{I} & \mathbf{K}_*^{\mathrm{T}} \\ \mathbf{K}_* & \mathbf{K}_{**} \end{bmatrix}\right). \tag{2}$$

Using the Gaussian conditioning rule, the predictive distribution becomes

$$\begin{aligned} \mathbf{f}_*|\mathbf{y} &\sim \mathrm{N}(\mathbf{f}_*|\boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*), \\ \boldsymbol{\mu}_* &= \mathbf{K}_*(\mathbf{K} + \sigma^2\mathbf{I})^{-1}\mathbf{y}, \\ \boldsymbol{\Sigma}_* &= \mathbf{K}_{**} - \mathbf{K}_*(\mathbf{K} + \sigma^2\mathbf{I})^{-1}\mathbf{K}_*^{\mathrm{T}}. \end{aligned} \tag{3}$$

Since the differentiation is a linear operator, the partial derivative of a Gaussian process remains a Gaussian process (Solak et al., 2003; Rasmussen and Williams, 2006). Thus, using partial derivative values for prediction and making predictions about the partial derivatives at a given point is easy. Since

$$\begin{aligned} \mathrm{cov}\left(\frac{\partial f^{(i)}}{\partial x_g^{(i)}}, f^{(j)}\right) &= \frac{\partial}{\partial x_g^{(i)}} \mathrm{cov}\left(f^{(i)}, f^{(j)}\right), \\ \mathrm{cov}\left(\frac{\partial f^{(i)}}{\partial x_g^{(i)}}, \frac{\partial f^{(j)}}{\partial x_h^{(j)}}\right) &= \frac{\partial^2}{\partial x_g^{(i)} \partial x_h^{(j)}} \mathrm{cov}\left(f^{(i)}, f^{(j)}\right) \end{aligned}$$

covariance matrices in Equations (1) and (2) can be extended to include partial derivatives either as observations or as values to be predicted.

Following Riihimäki and Vehtari (2010), denote by m the partial derivative value in the dimension j at $\tilde{\mathbf{x}}$. Then the probability of partial derivative is approximated using probit likelihood with a control parameter ν (see details in Riihimäki and Vehtari, 2010)

$$p\left(m \middle| \frac{\partial \tilde{f}}{\partial \tilde{x}_j}\right) = \mathbf{\Phi}\left(\frac{\partial \tilde{f}}{\partial \tilde{x}_j}\frac{1}{\nu}\right), \tag{4}$$

where

$$\mathbf{\Phi}(z) = \int_{-\infty}^{z} N(t \mid 0, 1)\mathrm{d}t.$$

Let $\mathbf{m}$ be a vector of q partial derivative values at $\tilde{\mathbf{X}} = \left(\tilde{\mathbf{x}}^{(1)}, \ldots, \tilde{\mathbf{x}}^{(q)}\right)$, $\mathbf{j}$ be a vector of the dimensions of the partial derivatives and $\tilde{\mathbf{f}}$ be the vector of latent values at $\tilde{\mathbf{X}}$. For shorthand, let the partial derivatives of latent values be $\tilde{\mathbf{f}}' = \left(\frac{\partial \tilde{\mathbf{f}}^{(1)}}{\partial \tilde{\mathbf{x}}^{(1)}_{\mathbf{j}^{(1)}}}, \ldots, \frac{\partial \tilde{\mathbf{f}}^{(q)}}{\partial \tilde{\mathbf{x}}^{(q)}_{\mathbf{j}^{(q)}}}\right)$. Assuming conditional independence given the latent derivative values, the likelihood becomes

$$p(\mathbf{m} \mid \tilde{\mathbf{f}}') = \prod_{i=1}^{q} \mathbf{\Phi}\left(\frac{\partial \tilde{\mathbf{f}}^{(i)}}{\partial \tilde{\mathbf{x}}^{(i)}_{\mathbf{j}^{(i)}}}\frac{1}{\nu}\right).$$

With function values at $\mathbf{X}$ and partial derivative values at $\tilde{\mathbf{X}}$, the joint prior for $\mathbf{f}$ and $\tilde{\mathbf{f}}$ then becomes

$$p\left(\begin{bmatrix}\mathbf{f}\\ \tilde{\mathbf{f}}'\end{bmatrix} \middle| \begin{bmatrix}\mathbf{X}\\ \tilde{\mathbf{X}}\end{bmatrix}\right) = \mathrm{N}\left(\begin{bmatrix}\mathbf{f}\\ \tilde{\mathbf{f}}'\end{bmatrix} \middle| \mathbf{0}, \begin{bmatrix}\mathbf{K}_{\mathbf{f},\mathbf{f}} & \mathbf{K}_{\mathbf{f},\tilde{\mathbf{f}}'}\\ \mathbf{K}_{\tilde{\mathbf{f}}',\mathbf{f}} & \mathbf{K}_{\tilde{\mathbf{f}}',\tilde{\mathbf{f}}'}\end{bmatrix}\right).$$

The joint posterior for latent values and latent value derivatives can be derived from the Bayes' rule

$$p(\mathbf{f}, \tilde{\mathbf{f}}' \mid \mathbf{y}, \mathbf{m}, \mathbf{X}, \tilde{\mathbf{X}}) = \frac{p(\mathbf{f}, \tilde{\mathbf{f}}' \mid \mathbf{X}, \tilde{\mathbf{X}})p(\mathbf{y} \mid \mathbf{f})p(\mathbf{m} \mid \tilde{\mathbf{f}}')}{Z}, \tag{5}$$

with the marginal likelihood

$$\begin{aligned} Z &= p(\mathbf{y}, \mathbf{m} \mid \mathbf{X}, \tilde{\mathbf{X}}) \\ &= \int p(\mathbf{f}, \tilde{\mathbf{f}}' | \mathbf{X}, \tilde{\mathbf{X}})p(\mathbf{y} \mid \mathbf{f})p(\mathbf{m} \mid \tilde{\mathbf{f}}')\mathrm{d}\mathbf{f}\mathrm{d}\tilde{\mathbf{f}}'. \end{aligned}$$

Since $p(\mathbf{m} \mid \tilde{\mathbf{f}}')$ is not Gaussian, the full posterior is analytically intractable and some approximation method must be used. Following Riihimäki and Vehtari (2010), we use expectation propagation for fast and accurate approximative inference.

Depending on the choice of covariance function, GPs have various amount of parameters that need to be selected. These hyperparameters include also the noise variance of the assumed Gaussian relationship between the observations and latent values. This topic of hyperparameter optimization is not covered here but the reader is further advised to see, for example, Rasmussen and Williams (2006) and Snoek et al. (2012).

2.2 BAYESIAN OPTIMIZATION WITH GAUSSIAN PROCESS PRIOR

Bayesian optimization is a sequential design strategy for global optimization of black-box functions (e.g., Močkus et al., 1978; Jones et al., 1998). The main steps of any BO algorithm are: First, the objective function is given a GP prior, which is updated with the selected evaluations so far. Second, an acquisition function, which depends on the posterior for the objective function, is used to decide what the next query point should be.

More explicitly, assume we have evaluated the objective function n times providing us the data $\mathbf{D} = \{y^{(i)}, \mathbf{x}^{(i)}\}_{i=1}^{n}$ where $y^{(i)}$ corresponds to the possibly noisy function evaluation at input location $\mathbf{x}^{(i)}$. We give a GP prior for the objective function and calculate its posterior distribution $y(\mathbf{x}') \mid \mathbf{D} \sim \mathrm{N}(\mu_*(\mathbf{x}'),\ \Sigma_*(\mathbf{x}'))$ where $\mu_*(\mathbf{x}')$ and $\Sigma_*(\mathbf{x}')$ are the predictive posterior mean and covariance functions of Equation (3). Furthermore, let the current best $\mu_{\text{min}} = \min_{1 \le i \le k}(\mu_*(\mathbf{x}^{(i)}))$. Then the three most established acquisition functions are:

1. *Expected improvement* (EI) by Jones et al. (1998), where the expected improvement is maximized over the current best:

$$a_{\text{EI}}(\mathbf{x}) = \mathrm{E}\left[\max(\mu_{\text{min}} - \mu_*(\mathbf{x}),\ 0)\right].$$

2. *Lower confidence bound* (LCB) by Srinivas et al. (2010), where the regret is minimized over the optimization area. Let η be a tuning parameter, then:

$$a_{\text{LCB}}(\mathbf{x}) = \mu_*(\mathbf{x}) - \sqrt{\eta^2 \Sigma_*(\mathbf{x})}.$$

   In our experiments, $\eta^2(t) = 2\log(t\frac{d}{2}+2\frac{\pi^2}{3\epsilon_\eta})$, with $\epsilon_\eta = 0.1$ Brochu E., Cora V. M., and de Freitas N. (2010).

3. *Probability of improvement* (PI) by Kushner (1964), where the next point is selected to be the one where probability of improving over the current best is the highest:

$$a_{\text{PI}}(\mathbf{x}) = \mathbf{\Phi}\left(\mu_{\text{min}} \mid \mu_*(\mathbf{x}),\ \Sigma_*(\mathbf{x})\right),$$

   with $\mathbf{\Phi}$ being a cumulative distribution function for a normal distribution.

3 BAYESIAN OPTIMIZATION WITH VIRTUAL DERIVATIVE SIGN OBSERVATIONS

When using Bayesian optimization for finding the minimum of a black box function, there usually is some prior

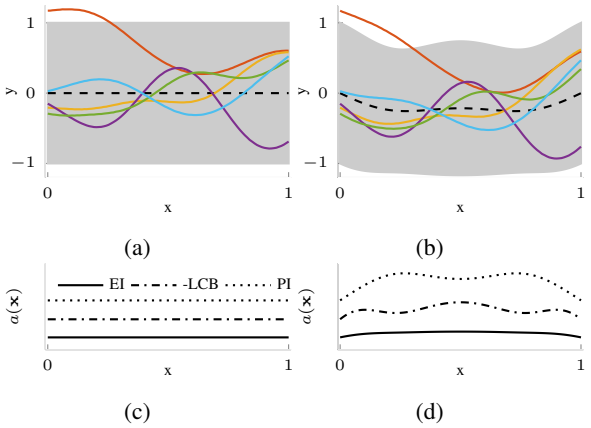

Figure 1: The GP prior in BO visualized (a) without virtual derivative observations, (b) with virtual derivative observations on both borders. Black dotted line is the mean of the GP, the light gray area is 68% central posterior interval and the five lines represent random function samples drawn from the prior. The acquisitions function values as a function of x are visualized for GPs (c) without virtual derivative observations, (d) with virtual derivative observations on both borders, for all three acquisition functions EI, LCB and PI. Since the next acquisition is chosen as maximum of EI and PI, but as a minimum of LCB, LCB is plotted as -LCB in order for the curves to be easily comparable.

knowledge available. First, the assumed smoothness of the function is used to select the covariance function and prior distributions for its hyperparameters. Second, some prior knowledge is needed to select the search bounds so that the minimum lies within them. Often the prior information for the bounds is vague and the global minimum or any of the local minima seldom lie exactly on any border. In other words, the gradient of the underlying function often points inwards on all the borders. To avoid the BO algorithm from taking samples from the border, virtual derivative observations can be added to the edges of the search space.

To encode the prior knowledge of the non-border minimum to the optimization algorithm, we propose the following dynamic approach. Just like in the regular BO with GP prior presented in the section 2.2, the objective function is given a GP prior which is updated according to the objective function evaluations so far. The next evaluation point is the acquisition function maximum, but if it is closer than threshold ϵ_b to the border of the search space, a virtual derivative observation is placed on the point that locates on the border of the search space and has smallest distance to the just proposed point. After having added this virtual observation, the GP prior is updated and new proposal for the next acquisition is computed. Algorithm 1 contains pseudo code for the proposed method.

Algorithm 1 Pseudo code of the proposed BO method. The acquisition methods a, the *stopping criterion* and the GP model nedd to be selcted beforehand.

while *stopping criterion* is False **do**
  Fit GP to the available dataset $\mathbf{X}, \mathbf{y}$.
  Optimise acquisition function, a, to find select new location $\mathbf{x}$ to evaluate.
  if $\mathbf{x}$ is close to the edge **then**
    Augment $\mathbf{X}$ with a virtual derivative sign observation at $\tilde{\mathbf{x}}$.
  else
    Augment $\mathbf{X}$ with $\mathbf{x}$ and evaluate g at $\mathbf{x}$.
  end if
end while

Since there is no actual knowledge about the true value of the partial derivatives on the borders, posterior distribution should behave robustly to deviations in the actual partial derivative values. By using ν very close to zero in Equation (4), the probit likelihood becomes very steep. This means that the likelihood values are close to 1 for all partial derivative values $m > 0$ and close to zero for all $m < 0$. Thus the assumed numeric partial derivative values do not matter and for instance $m_{d_i}^{(i)} = \pm 1$ can be used to express prior information on the borders. The effect of adding a virtual derivative observations on the borders of one dimensional function is visualized in the Figure 1. From the Figure it can be seen that the virtual derivative observations alter both the GP prior and the acquisition functions to resemble our prior belief of the location of the minimum.

Another parameter to be chosen is the threshold ϵ_b. As acquisitions closer than threshold are always rejected, ϵ_b should not be too large. Another argument to avoid too large values is the fading information value of the virtual observations. If the GP allows rapid changes in the latent values, virtual derivative observations affect to the posterior distribution only very locally and the added observations might affect to the posterior distribution only very marginally. However, if the threshold is too small, virtual observations are added very seldom and the algorithm differs only a little from the standard BO. Let l be the diameter of the search space. Our experiments suggest that $\epsilon_b \approx 0.01 \cdot l$ provides a good trade-off for adaptivity.

Basic time complexity of GP is $\mathcal{O}\left((n+q)^3\right)$, where n is the number of training samples and q is the number of partial derivative observations. If the search space is restricted to be a hyper-cube, partial derivative only in one direction is needed when a virtual observation is added on the edge of the search space, which helps to keep the

number of partial derivative observations relatively small. Naturally various sparse approximations could be used to improve the scaling of the GP computation, but here we focus on problems where we assume that acquisitions are costly and n is moderate. Thus, the time complexity of the algorithm itself is negligible compared to the time of acquisitions.

4 EXPERIMENTS

In this section we introduce the three case studies performed to gain insight about the performance of the proposed method. First, we introduce five different metrics for comparing performances of the algorithms, after this details of the experiments and results are presented.

4.1 METRICS FOR COMPARING PERFORMANCES

There are no standard methods for comparing performances of different global optimization approaches. We propose five metrics which we think provide a good overview of the differences in performances. In the following, GPD is the GP with virtual derivative observations, GP is a standard GP, d is the dimensionality of the search space and, n is the total number of acquisitions obtained. In addition, let $g(\mathbf{x})$ be the noiseless black box function to be optimized, g_{min} be the absolute minimum of that function, $\mathbf{x}_{\text{min}}$ be the point where that minimum is located and let acquisitions obtained by using two different methods be denoted by $\{\mathbf{x}_{\text{GP}}^{(i)}\}_{i=1}^{n}$ and $\{\mathbf{x}_{\text{GPD}}^{(i)}\}_{i=1}^{n}$. Furthermore, let

$$g_{\text{ref min}} = \max\left(\min_{1\leq i\leq n}\left(g(\mathbf{x}_{GPD}^{(i)})\right), \min_{1\leq i\leq n}\left(g(\mathbf{x}_{GP}^{(i)})\right)\right).$$

be the maximum of the absolute minimums that the compared algorithms are able to find,
$h_{\text{GPD}} = \min_{1\leq i\leq n}(i \,|\, g(\mathbf{x}_{\text{GPD}}^{(i)}) \leq g_{\text{ref min}})$ and
$h_{\text{GP}} = \min_{1\leq i\leq n}(i \,|\, g(\mathbf{x}_{\text{GP}}^{(i)}) \leq g_{\text{ref min}})$ be the first indices when the acquisitions are better or equal to $g_{\text{ref min}}$. Then the five metrics are:

- *Percentual minimum difference* (PMD): Let $\mathbf{x}_{GPD}^{(i)}$ and $\mathbf{x}_{GP}^{(j)}$ be the acquisitions that minimize the GP mean functions $\mu_{GPD}(\mathbf{x})$ and $\mu_{GP}(\mathbf{x})$, then:

  $$\text{PMD} = \frac{g(\mathbf{x}_{\text{GPD}}^{(i)}) - g(\mathbf{x}_{\text{GP}}^{(j)})}{\max(g(\mathbf{x}_{\text{GPD}}^{(i)}) - g_{\text{min}},\, g(\mathbf{x}_{\text{GP}}^{(j)}) - g_{\text{min}})}.$$

  PMD is designed to compare the absolute performances of the algorithms and intuitively it measures the difference of the best values of both algorithms. *The interpretation for the magnitude of the value is that negative values tell that the proposed method is better, the values are always scaled between* -1 *and* 1 *and the further away the value is from 0, the bigger the difference between the two methods is.*

- *Percentual hit difference* (PHD):

  $$\text{PHD} = \frac{h_{\text{GPD}} - h_{\text{GP}}}{\max(h_{\text{GPD}}, h_{\text{GP}})}$$

  PHD is created for comparing the speeds of the algorithms and intuitively it measures difference of how fast both algorithms are able to find good enough values. The interpretation of the magnitude of the value is exactly the same as for PMD.

- *Percentual border hit difference* (PBHD): Let b_{GPD} be the number of evaluations of GPD and b_{GP} be the number of evaluations of GP that are within ϵ_{BHD} from any border of the search space. Then

  $$\text{PBHD} = \frac{b_{\text{GPD}} - b_{\text{GP}}}{\max(b_{\text{GPD}}, b_{\text{GP}})}.$$

  Assuming that the minimum is not near the border, BHD tells the scaled difference of unnecessary samples taken near the borders. The interpretation for the magnitude of the value is exactly the same as for PMD. In our experiments, $\epsilon_{\text{BHD}} = 0.05$.

- *Average evaluation distance difference* (AED): Let $\text{dist}(\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$ be Euclidean distance, then $d_{\text{GPD}} = \sum_{i=1}^{h_{\text{GPD}}} \frac{\text{dist}(\mathbf{x}_{\text{GPD}}^{(i)}, \mathbf{x}^{\text{min}})}{h_{\text{GPD}}}$ and $d_{\text{GP}} = \sum_{i=1}^{h_{\text{GP}}} \frac{\text{dist}(\mathbf{x}_{\text{GP}}^{(i)}, \mathbf{x}^{\text{min}})}{h_{\text{GP}}}$ are the average distances of acquisitions from the true minimum before finding $g_{\text{ref min}}$. Then:

  $$\text{AED} = \frac{d_{\text{GPD}} - d_{\text{GP}}}{\max(d_{\text{GPD}},\, d_{\text{GP}})},$$

  Intuitively, AED measures the overall performance of the algorithm before finding the minimum. The interpretation for the magnitude of the value is exactly the same as for PMD.

- *Virtual derivative observations per dimension* (VDO): Let q be the number of d dimensional virtual derivative observations, then

  $$\text{VDO} = \frac{q}{d}.$$

  Unlike other metrics, VDO doesn't compare the two methods, but give hint about how much the computational complexity of the proposed method differs from the standard BO with GP prior. Intuitively, larger VDOs are worse, since they increase the computational burden of the algorithm as GP's scale as $\mathcal{O}\left((n+q)^3\right)$.

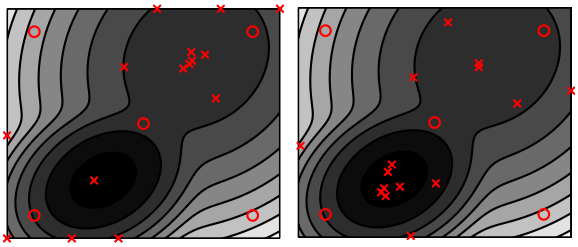

(a) (b)

Figure 2: 15 acquisitions with (a) standard BO (b) BO with virtual derivative observations. In both figures the five red circles are the points used to initialise the GP and the 15 red crosses are the acquisitions. In both algorithms, the used acquisition function is LCB. In both figures, darker colors represent lower function values.

4.2 EXPERIMENTAL SETUP

The proposed Bayesian optimization algorithm was implemented in GPstuff toolbox (Vanhatalo et al., 2013). In both case studies to be covered later, we use zero mean GP prior for regular observations and probit likelihood with scaling parameter $\nu = 10^{-9}$ for the virtual derivative observations. In addition to this, we use the squared exponential covariance function with each dimension having their own length scales,

$$k(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left\{-\frac{1}{2}\sum_{i=1}^{d}\frac{(\mathbf{x}_i - \mathbf{x}_i')^2}{l_i^2}\right\},$$

where d is the input space size, σ_f^2 is magnitude and l_i are characteristic length scales in each dimension. The partial derivatives are

$$\begin{aligned}
\frac{\partial k(\mathbf{x}, \mathbf{x}')}{\partial \mathbf{x}_j} &= -\sigma_f^2 \exp\left\{-\frac{1}{2}\sum_{i=1}^{d}\frac{(\mathbf{x}_i - \mathbf{x}_i')^2}{l_i^2}\right\} \\
&\quad \times \frac{\mathbf{x}_j - \mathbf{x}_j'}{l_j^2}, \\
\frac{\partial k(\mathbf{x}, \mathbf{x}')}{\partial \mathbf{x}_j'} &= -\frac{\partial k(\mathbf{x}, \mathbf{x}')}{\partial \mathbf{x}_j}, \\
\frac{\partial k(\mathbf{x}, \mathbf{x}')}{\partial \mathbf{x}_j \mathbf{x}_k'} &= \sigma_f^2 \exp\left\{-\frac{1}{2}\sum_{i=1}^{d}\frac{(\mathbf{x}_i - \mathbf{x}_i')^2}{l_i^2}\right\} \\
&\quad \times \left(-\frac{\mathbf{x}_j - \mathbf{x}'_j}{l_j^2}\frac{\mathbf{x}_k - \mathbf{x}_k'}{l_k^2} + \frac{\delta_{jk}}{l_j^2}\right),
\end{aligned}$$

which are needed to compute the virtual derivatives.

Initial acquisitions are generated with fractional factorial design (Croarkin et al., 2013). In the experiments, we use the minimum amount of initial points so that they still fill the input space (see table 3.17 from the citation).

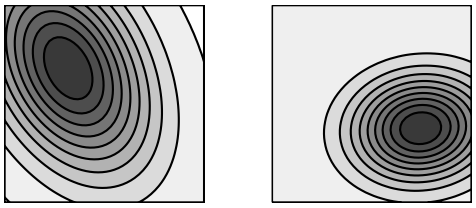

Figure 3: Two examples of two dimensional random normal distribution functions. In both examples, the figure bounds represent the search space, darker areas represent lower function values.

For the presented BO algorithm, virtual derivative observations are added if the next proposed point is within 1% of the length of the edge of the search space to any border.

4.3 Case study 1: A simple example function

The algorithm is used to illustrate the unwanted boundary over-exploration effect of the regular BO. To show this, simple function consisting of two Gaussian components is optimized with the standard BO and BO with virtual derivative observations using LCB as an acquisition function. The function and 15 first acquisitions are visualized in Figure 2.

The results show that regular BO has sampled seven, and the proposed method has sampled four points from near the borders. Furthermore, regular BO has only one acquisition from near the true minimum, whereas the proposed method has eight acquisitions.

4.4 Case study 2: Random multivariate normal distribution functions

The algorithm is used to find the minimum of different d-dimensional multivariate normal distribution (MND) functions:

$$g(\mathbf{x}) = -\exp\left\{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^{\mathrm{T}}\boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right\},$$

where the means $\boldsymbol{\mu}$ are selected uniformly at random such that $0.2 \leq \boldsymbol{\mu}_i \leq 0.8 \;\forall\; i = 1, \ldots, d$, covariances $\boldsymbol{\Sigma}$ are a random positive semidefinite matrices which eigenvalues $\mathbf{e}$ are selected uniformly at random such that $\frac{1}{70} \leq \mathbf{e}_i \leq \frac{1}{7} \;\forall\; i = 1, \ldots, d$. To simulate real life observations, random noise $\epsilon \sim N(0, s)$ is added to the observations $y(\mathbf{x}) = g(\mathbf{x}) + \epsilon$, for some fixed s. Two example functions, without noise, are visualized in Figure 3.

We drew 100 MND functions in each dimension as $d = 1, \ldots, 11$, for three different noises $s = \{0, 0.05, 0.1\}$. For all these 3300 functions, we ran 100 iterations for thetwo BO algorithms for all three acquisition functions,

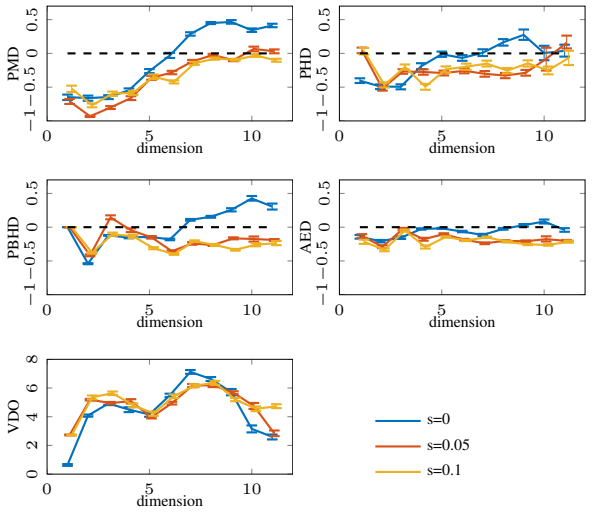

Figure 4: Means of different metrics as a function of the dimension of the random multivariate normal distribution functions. Values under zero represent better performance of the GPD method with respect to the baseline. Different continuous lines resemble different amounts of added noise to the black box function evaluations. Error bars around the continuous lines are 25th and 75th percentiles of the means. The acquisition function of both compared methods is EI. Every point in every plot is computed as a mean from 100 separate algorithm runs for two different algorithms.

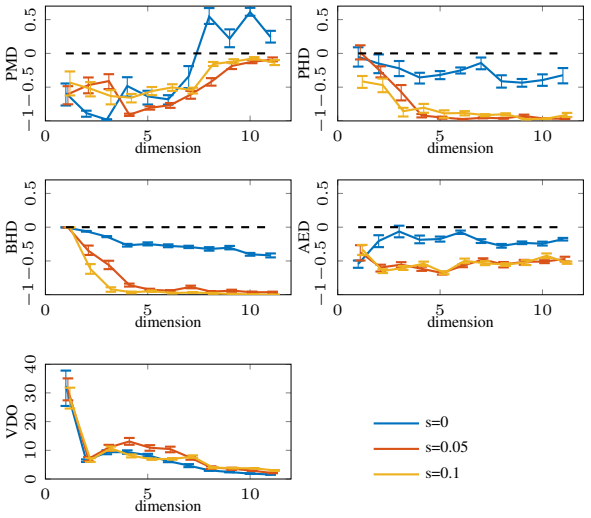

Figure 5: Same as in Figure 4, but for LCB as an acquisition function.

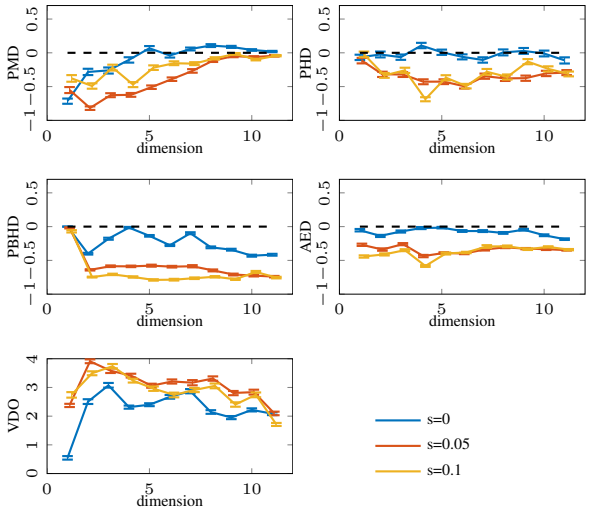

Figure 6: Same as in Figure 4, but for PI as an acquisition function.

EI, LCB and PI. The means and 25th and 75th percentiles for those runs for all five metrics described in the section 4.1 are visualized in Figures 4, 5 and 6 for acquisition functions EI, LCB and PI. The plotted quantiles are obtained using Bayesian bootstrap (Rubin, 1981). In addition to the metrics after 100 iterations, Figure 7 visualizes the development of the means of the metrics as a function of number of iterations for EI as ans acquisition function for functions of dimensions 2, 6 and 10 with $s = 0.05$. The development is shown only for one acquisition function and only in few dimensions as the results are very similar for all settings.

The results show that in most dimensions for most noise levels and for all acquisition functions, the proposed method performs better than the baseline method on average. Particularly, the results for PMD show that the proposed method is able to find better minima than the baseline, especially on lower dimensions. Furthermore, the results for PHD show that in all dimensions, and especially for noisy functions, the proposed method is able to find smaller values significantly faster. In addition to this, PBHD shows, as expected, that the proposed method samples less near the borders, and AED shows that the samples are taken closer the true minimum. The combining factor for all these metrics is that the improvement is more significant if the observation are noisy. Furthermore, the only cases where the proposed methods performance is worse is in high dimensions and for the noiseless function observations. The reason for this stays unknown, but we suspect that in these cases the error of EP approximation outweighs the benefit of added partial derivative observations and that the results might be different if, for instance, MCMC was used to generate the posterior distribution of the GP.

In addition to the results after 100 acquisitions, as seen from Figure 7, the proposed method seems to be relatively robust to the number of iterations taken during BO. Af-

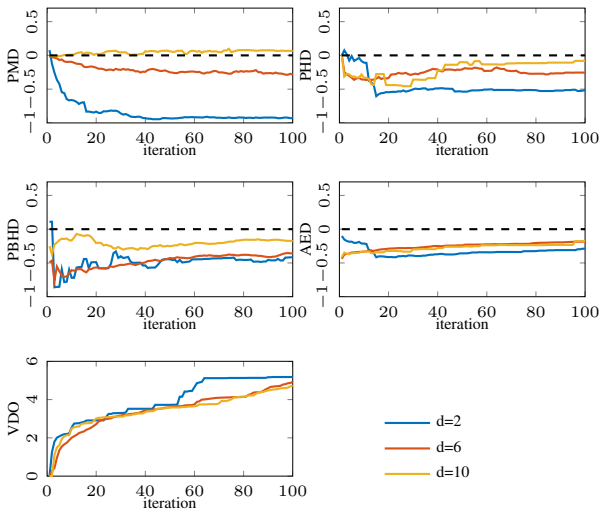

Figure 7: Means of different metrics as a function of optimization length for functions of different dimensions. The used functions are 2, 6 and 10 dimensional multivariate normal distribution functions with $s = 0.05$. The acquisition function of both compared methods is EI. *The values after 100 iteration correspond to the values of the corresponding dimension of the red line in Figure 4.*

ter the first 20 iterations, all metrics change slowly for functions of all dimensions.

4.5 Case study 3: Sigopt function library

Dewancker et al. (2016) introduced a benchmark function library[1] called Sigopt for evaluating BO algorithms. This library consists of functions of varying dimensionality and each of which belonging to one or more groups like unimodal, discrete or oscillatory functions. In addition to the functions, the library also defines the search space and true minimum of these functions. When restricting the functions to at maximum 11 dimension and taking into account only unimodal functions with no global border minima, the library outputs 113 functions. For these functions, the number of functions per dimension varies between 8 and 21. As in the previous case study, to mimic real use cases, the function observations are corrupted with added random noise $y(\mathbf{x}) \sim g(\mathbf{x}) + \mathrm{N}(0, s^2)$, for $s = \{0, 0.05, 0.1\}$. Two example functions, without noise, are visualized in Figure 8.

Similarly as in the previous example, for all these functions from Sigopt dataset, we ran 100 iterations for the proposed and baseline methods for all acquisition functions. The means and 25th and 75th percentiles for those runs for all five metrics described in the section 4.1 are

[1]Function library available at: https://github.com/sigopt/evalset

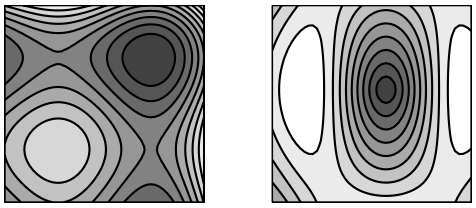

Figure 8: Two examples of two dimensional functions from the Sigopt function library. In both examples, the plot bounds represent the search space, darker areas represent lower function values.

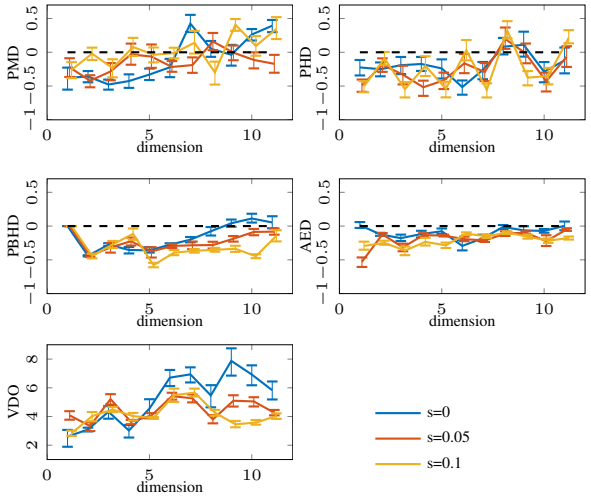

Figure 9: Means of different metrics as a function of the dimension of the input space of functions of Sigopt dataset. Values under zero represent better performance of the GPD method with respect to the baseline. Different continuous lines resemble different amounts of added noise to the black box function evaluations. Error bars around the continuous lines are 25th and 75th percentiles of the means. The acquisition function of both compared methods is EI.

visualized in Figures 9, 10 and 11 for acquisition functions EI, LCB and PI. As previously, the plotted quantiles are obtained using Bayesian bootstrap. In addition to the metrics after 100 iterations, Figure 12 visualizes the development of the means of the metrics as a function of number of iterations for EI as acquisition function for functions of dimensions 2, 6 and 10 with $s = 0.05$. The development is shown only for one acquisition function and only in few dimensions as they are very similar for all settings.

The results are very similar to the results of random MND functions. However, on average, the results are not as good, and the deviation of results between different functions is bigger. The reason for results being worse is explained by the fact that for some functions, there are

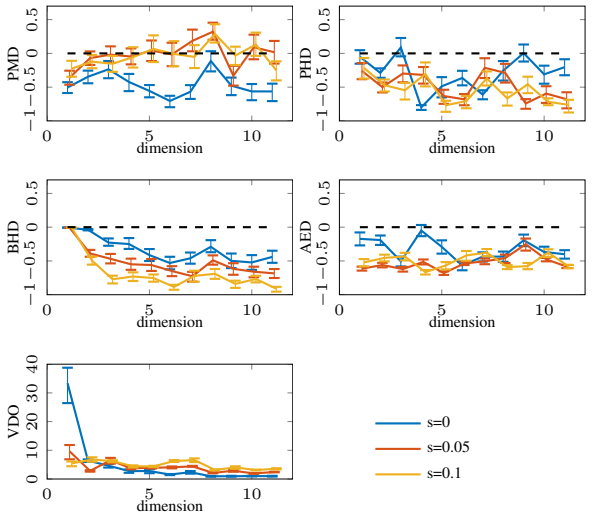

Figure 10: Same as in Figure 9, but for LCB as an acquisition function.

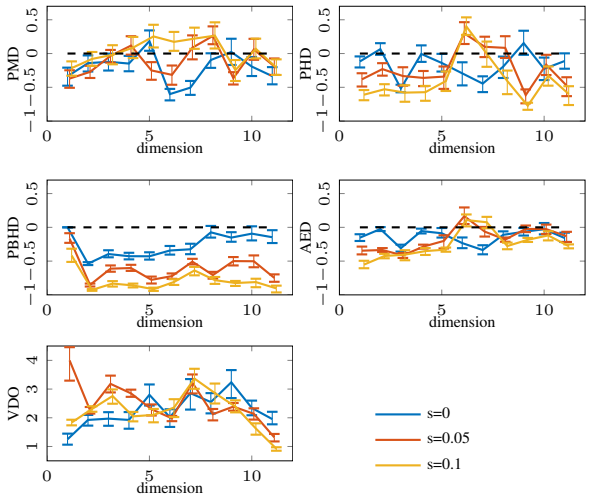

Figure 11: Same as in Figure 9, but for PI as an acquisition function.

local minima on borders (see examples from Figure 8). Furthermore, the larger deviation is explained by fewer functions in each dimension. This also explains the larger variation in the the progressions of different metrics of Figure 12. However, even though the results are not as good as previously, the proposed method still performs better than the baseline method on average, and so it can be sued as a default for BO.

5 CONCLUSIONS

We have presented here a Bayesian optimization algorithm which utilizes qualitative prior information concerning the objective function on the borders. Namely,

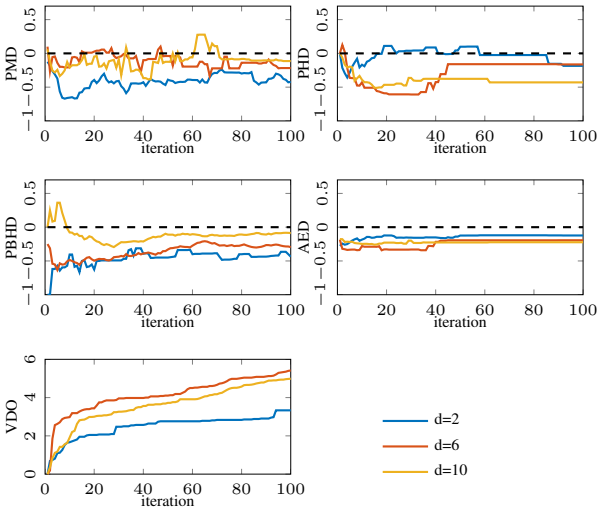

Figure 12: Means of different metrics as a function of optimization length for functions of different dimensions. The used functions are 2, 6 and 10 dimensional multivariate normal distribution functions with $s = 0.05$. The acquisition function of both compared methods is EI.*The values after 100 iteration correspond to the values of the corresponding dimension of the red line in Figure 9.*

we assume that the gradient of the underlying function points towards the centre on all borders. Typical uses of Bayesian optimization concern expensive functions and in many applications qualitative knowledge of the generic properties of the function are known prior to optimization. Our approach generalizes naturally to these settings so that the user could inform the GP prior about the direction of gradient at any location of the search space.

The proposed BO method has proved to significantly improve the optimization speed and the found minimum when comparing the average performance to the performance of the standard BO algorithm without virtual derivative sign observations. The difference in performance is more significant if the assumption of non-existent global or local minima on the border of the search space holds, but is still notable if the assumption is relaxed so that the global minimum is not located on the border.

Acknowledgments

We thank Juho Piironen for his valuable comments to improve the manuscript. We also acknowledge the computational resources provided by Aalto University's Department of Computer Science.

References

Brochu E., Cora V. M., and de Freitas N. (2010). A

tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*. 3

Croarkin, C., Tobias, P., and Zey, C. (2013). *Engineering statistics handbook*. NIST iTL. 6

Dewancker, I., McCourt, M., Clark, S., Hayes, P., Johnson, A., and Ke, G. (2016). A stratified analysis of Bayesian optimization methods. *arXiv preprint arXiv:1603.09441*. 8

González, J., Osborne, M. A., and Lawrence, N. D. (2016). GLASSES: relieving the myopia of Bayesian optimisation. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS 2016, Cadiz, Spain, May 9-11, 2016*, pages 790–799. 1

Gosling, J. P., Oakley, J. E., and O'Hagan, A. (2007). Nonparametric elicitation for heavy-tailed prior distributions. *Bayesian Analysis*, 2(4):693–718. 2

Jones, D. R., Schonlau, M., and Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492. 3

Koistinen, O., Maras, E., Vehtari, A., and Jónsson, H. (2016). Minimum energy path calculations with Gaussian process regression. *Nanosystems: Physics, Chemistry, Mathematics*, 7(6):925–935. 2

Krause, A. and Guestrin, C. (2007). Nonmyopic active learning of Gaussian processes: an exploration-exploitation approach. In *Proceedings of the 24th International Conference on Machine Learning*, volume 227 of *ACM International Conference Proceeding Series*, pages 449–456. ACM. 1

Kushner, H. J. (1964). A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering*, 86(1):97–106. 3

Michael, J. and Víctor, P. (2016). Bayesian optimization with shape constraints. *arXiv preprint arXiv:1612.08915*. 2

Močkus, J., Tiesis, V., and Žilinskas, A. (1978). The application of Bayesian methods for seeking the extremum. In Dixon, L. and Szego, G., editors, *Toward Global Optimization*. 2nd edition, Elsevier. 3

O'Hagan, A. (1992). Some Bayesian numerical analysis. In Bernardo, J. M., Berger, J. O., Dawid, A. P., and Smith, A. F. M., editors, *Bayesian statistics 4*, pages 345–363. Oxford University Press. 2

Osborne, M. A., Garnett, R., and Roberts, S. J. (2009). Gaussian processes for global optimization. In *3rd international conference on learning and intelligent optimization (LION3)*, pages 1–15. 1

Rasmussen, C. E. and Williams, C. K. I. (2006). Gaussian processes for machine learning. *The MIT Press*, 2(3):4. 2, 3

Riihimäki, J. and Vehtari, A. (2010). Gaussian processes with monotonicity information. In *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*, pages 645–652. 2, 3

Riihimäki, J. and Vehtari, A. (2014). Laplace approximation for logistic Gaussian process density estimation and regression. *Bayesian Analysis*, 9(2):425–448. 2

Rubin, D. B. (1981). The Bayesian bootstrap. *The Annals of Statistics*, 9(1):130–134. 7

Shahriari, B., Bouchard-Côté, A., and de Freitas, N. (2016a). Unbounded Bayesian optimization via regularization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, pages 1168–1176. 1

Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and de Freitas, N. (2016b). Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175. 1

Siivola, E., Piironen, J., and Vehtari, A. (2016). Automatic monotonicity detection for Gaussian processes. *arXiv preprint arXiv:1610.05440*. 2

Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems*, pages 2951–2959. 3

Solak, E., Murray, S. R., Leithead, W. E., Leith, D. J., and Rasmussen, C. E. (2003). Derivative observations in Gaussian process models of dynamic systems. In *Advances in Neural Information Processing Systems*, pages 1033–1040. 2

Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. (2010). Gaussian process optimization in the bandit setting: No regret and experimental design. *Proceedings of the 27th International Conference on Machine Learning*. 3

Vanhatalo, J., Riihimäki, J., Hartikainen, J., Jylänki, P., Tolvanen, V., and Vehtari, A. (2013). GPstuff: Bayesian modeling with Gaussian processes. *Journal of Machine Learning Research*, 14:1175–1179. 6

Wang, X. and Berger, J. O. (2016). Estimating shape constrained functions using Gaussian processes. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):1–25. 2

Wu, J., Poloczek, M., Wilson, A. G., and Frazier, P. I. (2017). Bayesian optimization with gradients. *arXiv preprint arXiv:1703.04389*. 2